

Statement Of Purpose - Stanford CS

Mitchell Gordon (m.gordon@rochester.edu)

My objective is to pursue a Ph.D. in Computer Science with a focus on technical human-computer interaction. My past research in HCI has investigated: (a) how to use coordinated crowds of people to create intelligent, low-latency systems for data analysis [2, 3, 4, 5], (b) using volunteer crowds of learners to generate new knowledge that previously required tedious expert labor [1], and (c) novel interfaces for intelligent text input on smartwatch devices [6].

My motivation to pursue a Ph.D comes from my experience working on several substantial research projects as an undergraduate. I've worked in groups at Google (Google Keyboard), CMU (HCII), and the University of Rochester and have also been a software engineering intern at Google (YouTube). Through these experiences, I have found that I am most engaged when working on ambitious technical research projects. I am interested in projects that, if successful, have the potential to impact a large number of people by improving the state of the art in a given area. But it is equally important to me that I generalize the new knowledge that I create so it can be used by other researchers in their own projects. Ultimately, my goal is to be a professor at a research university.

My interest in HCI research first developed while working with Professors Walter Lasecki and Jeffrey Bigham on several crowd-powered systems. One of our most interesting projects is Glance, a system that I built with Professors Lasecki and Bigham with further advising from Professor Steven Dow. Glance uses a large coordinated crowd of parallel online workers to perform video analysis ("coding"). In a second-author paper at UIST, we found that we could reduce the time it takes to perform video coding from hours to minutes, making it possible for researchers to gain new insights into their data faster than before [3]. We then investigated the privacy implications of using a crowd to analyze videos, with the goal of identifying methods to obscure identity while retaining awareness of actions in video. In a second-author paper at CHI, we used sample videos with filters applied and measured how accurately the crowd could identify subjects versus how accurately they could code the videos. We found that we could code for many key behaviors in video without compromising participant identity [2].

I also showed that Glance can be used to help evaluate the accessibility of websites in a paper that won the Student Research Competition (SRC) at ASSETS [4] and won 2nd place at the ACM SRC Grand Finals, a competition between the SRC winners of ACM conferences in all subfields of Computer Science, and I was recognized for this prize at the ACM Awards Banquet. The Glance system has been received with a great amount of interest by researchers in the social sciences at Stanford, CMU, Michigan, and several other institutions. To facilitate Glance and other projects, I built LegionTools [5], a toolkit and UI for Amazon Mechanical Turk that makes it possible to recruit and manage workers for multiple large real-time experiments. LegionTools is currently being used by research groups at CMU, Michigan, and Rochester.

After working on several crowdsourcing projects, I was excited to apply what I had learned to work on a problem I care about deeply: CS education. Last spring I led the Codepourri project under the supervision of Prof. Philip Guo. Unlike in math or other sciences, step-by-step tutorials for computer programming have previously been difficult and tedious to create, as they required making a static representation of the dynamic state in which a wide variety of possible steps can occur. Codepourri solves this problem; it enables people to easily create visual coding tutorials by annotating steps in an automatically generated program visualization, and uses a volunteer crowd of asynchronous learners to generate the annotations. Using a crowd of learners is both more scalable than a crowd of experts and allows us to overcome the expert blind spot, the phenomenon in which experts forget the experience of being a novice and skip over rudimentary explanations that novices need to see. Our evaluation found that Codepourri could create better tutorials than experts, and I published Codepourri as a first-author full paper at VL/HCC [1].

Most recently I was a research intern at Google working with Dr. Shumin Zhai on the Google Keyboard team. I led a research project that addressed the problem of creating an official Android Wear on-screen keyboard for smartwatches. Keyboards did not previously exist for either Android Wear or Apple's iWatch outside of the research literature. Given Google confidentiality, I am unable to provide more detail at this time, however my findings revealed some surprising conclusions which overturn the previous

direction of research in this area. Further, they were instrumental in making fundamental design and technical decisions for a product that will be used by millions of people. I submitted a first-author full paper to CHI on this work which is currently under review [6].

Though I am open to a wide variety of research within HCI, my experience building crowd-powered systems and developing educational technologies has inspired an interest in intelligent programming assistants that make use of online learners. I intend to harness the power of the work that is already being done by people learning to code to eliminate some of the tedium and verbosity that is currently necessary for writing commonly used logic in code. More concretely, I am interested in leveraging my previous work to build IDEs that can make existing computer programmers become more productive by eliminating some of the duplication of effort that occurs among millions of coders. Aside from this research, I am always open to interesting technical projects in other areas of HCI, and hope to gain a deep understanding of the field during my Ph.D.

At Stanford, there are a few professors whose projects are especially interesting to me: Michael Bernstein, Maneesh Agrawala, and James Landay . After reading several papers in each of these groups, I see a clear fit for my skills and interests at Stanford and am confident that it is a great place for me to pursue a Ph.D.

- [1] **Mitchell Gordon**, Philip J. Guo. "Codepourri: Creating Visual Coding Tutorials Using A Volunteer Crowd Of Learners." in VL/HCC 2015. Atlanta, GA.
- [2] Walter S. Lasecki, **Mitchell Gordon**, Winnie Leung, Ellen Lim, Jeffrey P. Bigham, Steven P. Dow. "Exploring Privacy and Accuracy Trade-Offs in Crowdsourced Behavioral Video Coding." in CHI 2015. Seoul, Korea.
- [3] Walter S. Lasecki, **Mitchell Gordon**, Malte Jung, Danai Koutra, Steven P. Dow, Jeffrey P. Bigham. "Glance: Rapidly Coding Behavioral Video with the Crowd." in UIST 2014. Honolulu, HI.
- [4] **Mitchell Gordon**. "Web Accessibility Evaluation with the Crowd: Using Glance to Rapidly Code User Testing Video." ASSETS 2014, Student Research Competition. Rochester, NY.
- [5] **Mitchell Gordon**, Jeffery P. Bigham, Walter S. Lasecki. "LegionTools: A Toolkit + UI for Recruiting and Routing Crowds to Synchronous Real-Time Tasks." UIST 2015. Charlotte, NC.
- [6] **Mitchell Gordon**, Tom Ouyang, Shumin Zhai. "WatchWriter: Tap and Gesture Typing on a Smartwatch with Miniature Qwerty and Beyond." in CHI 2016. **Currently Under Review**